

DRAFT Trinity/NERSC-8 Mini-application and SSP Benchmark Instructions

Introduction

Application benchmarks and microbenchmarks will play a critical role in evaluation of the offered system. The Trinity/NERSC-8 benchmarks serve three purposes:

1. The application benchmarks have been carefully chosen to represent characteristics of the expected Trinity and NERSC-8 workloads, both of which consist of solving complex scientific problems using diverse computational techniques at high degrees of parallelism.
2. The benchmarks give the Offeror the opportunity to provide NERSC and ACES with concrete data associated with the performance and scalability of the proposed system on applications that NERSC and ACES consider programmatically important.
3. The benchmarks will be used as an integral part of the system acceptance test and as a measurement of performance throughout the operational lifetime of the systems.

The Trinity/NERSC-8 benchmarks comprise tests at varying levels of the benchmark hierarchy that range from system component-level tests and kernels to Mini-Applications to Full Application Benchmarks. An aggregate performance measure, called the Sustained System Performance (SSP) metric, will be calculated based on application performance. The Offeror should pay particular attention to the SSP calculation, as it is one of the key metrics for system evaluation in this procurement.

Submission Guidelines

Benchmark results (or projections including original results) for the proposed system shall be recorded in the spreadsheet provided. Note that in the supplied spreadsheet, the "Proposed System Node Count" entry refers to the value for the full, proposed system, whether benchmarked or projected. Additionally, the Offeror shall submit electronically all completed tables, benchmark source codes, compile/build scripts, output files and documentation on any code optimizations or configuration changes on a CD or similar medium. The submitted source code shall be in a form that can be readily compiled on the proposed system. Do not include object and executable files, core dump files or large binary data files in the electronic submission. An audit trail showing any changes made to the benchmark codes must be supplied and it must be sufficient for ACES and NERSC to determine that the changes made conform to the spirit of the benchmark and do not violate any specific restrictions on the various benchmark codes.

If performance projections are used, this must be clearly indicated. The Trinity/NERSC-8 consortium will be the sole judge of the validity of any projected results. The output files on which the projections are based, and a description of the projection method must be included. In addition, each system used for benchmark projections must be described. Descriptions of the benchmarking system should be included in the supplied benchmarking spreadsheet.

Run Rules

Specific run rules for each benchmark will be included with the individual benchmark source code distribution, supplying specific requirements and instructions for compiling, executing, verifying numerical correctness and reporting results. Benchmark performance shall only be accepted from runs that exhibit correct execution. Only software tools and libraries that will be included for general use in the proposed system as supported product offerings are permissible for building and executing the benchmarks.

Benchmark Descriptions

Microbenchmark Tests

The Microbenchmark Tests, listed in Table 1, are simple, focused tests that are easily compiled and executed. The results allow a uniform comparison of features and provide an estimation of system balance. Descriptions and requirements for each test are included in the source distribution for each microbenchmark. The results for the proposed system shall be recorded in the provided benchmarking spreadsheet. Note: not all microbenchmarks need be submitted to respond to the RFP as several are only required at system acceptance. The stage at which a response is required for a specific microbenchmark is indicated in Table 1.

Modifications to the microbenchmarks are only permissible to enable porting and correct execution on the target platform.

Table 1. Lower Level Tests (Microbenchmarks)

Benchmark		Purpose	RFP Response	Acceptance Test
STREAM	Memory bandwidth		X	X
PSNAP	OS jitter			X
IOR	Sequential & parallel I/O performance		X	X
MDTEST	Filesystem metadata server performance		X	X

OMB	Interconnect performance	X	X
PYNAMIC	Dynamic loading and dynamic libraries		X
SMB	Message passing host processor overhead	X	X
ZIATEST	MPI application launch performance		X
UPC FT	PGAS functionality and performance		X
MPIMEMU	MPI node memory usage	X	X

Mini-Application Benchmarks

The application benchmarks are a representation of the ACES and NERSC workloads and span a variety of algorithmic and scientific spaces. The list of application benchmarks is shown in Table 2. Each application contains a README file describing how to build and run each application as well as any supporting library requirements.

For each application there are three problem sizes to be reported: a “small” problem to be run on a single node, a large problem that is weak-scaled to run on order-1,000 nodes, and an “extra large” weak-scaled problem that is run on order-10,000 nodes. These problems shall be at least eight times (8x) larger than a node’s last level cache (and scratchpad memory if applicable, e.g. GPGPU GDDR memory). It is desired that at least 50% of the proposed system’s aggregate main memory (i.e., 0.5x, 500x and 5,000x single node main memory capacity respectively) be utilized by the application. Given these guidelines, it may be necessary to modify an extra-large problem, in conjunction with the Trinity/NERSC-8 consortium, to fit a particular architecture.

For each problem size (small, large and extra large), two cases must be reported: an unoptimized “base case” that uses “as-is” code along with an MPI-only execution model and a “highly optimized” case that allows the Offeror broad latitude to optimize code and demonstrate the best-case performance potential of the system, especially for different execution models (MPI+X). ***It is extremely important for the Offeror to provide results for each benchmark.***

Table 2: Mini-Application Benchmarks

Application	Discipline	Small / Large	Extra Large
MiniFE	Finite elements	X	X
MiniGhost	Halo exchange	X	X
MiniContact	Geometric contact	X	X
UMT	Unstructured-mesh deterministic radiation Transport	X	X
AMG	Algebraic Multigrid	X	X
SNAP	Discrete Ordinates particle transport	X	X
GTC	Plasma physics	X	
MILC	Lattice QCD	X	
MiniDFT	Density functional theory	X	

Base Case

The base case requires an MPI-only execution model and limits the scope of optimization. This is necessary to provide a point of reference relative to known systems and to ensure that any proposed system can run legacy codes that use the MPI-only execution model. For the base case, any existing APIs in the codes that exploit additional parallelism, such as OpenMP, may not be enabled. The base case will be used to understand baseline application for applications currently using only MPI and will be used to understand the potential for application performance improvement when compared with the optimized case. In the base case, for all applications, modifications are permissible only to enable porting and correct execution on the target platform. No changes related to optimization are permissible. Library routines may be used as long as they currently exist in an Offeror's supported set of general or scientific libraries, and as long as they will be in such a set when the system is delivered. Any libraries used must not specialize or limit the applicability of the benchmark nor violate the measurement goals of the particular benchmark. Source preprocessors, execution profile feedback optimizers, etc., are allowed as long as they are, or will be, available and supported as part of the compilation system for the delivered systems. Only publicly available and documented compiler switches shall be used. Compiler optimizations will be allowed only if they do not increase the runtime or artificially increase the delivered floating-point operation rate per second by performing non-useful work.

The vendor will provide results for two benchmarks: a small single-node MPI-only case and a larger multi-node MPI-only case. In both cases, the vendor is free to choose the MPI concurrency and layout (e.g., affinity or hardware multithreading) that minimizes execution time. The rationale for these choices must be detailed in the response. Note that the number of MPI tasks that can be used for a particular benchmark may be constrained by domain decomposition rules inherent in the code.

Optimized MPI+X Cases

The optimized case allows the Offeror to highlight the features and benefits of the proposed system by submitting benchmarking results obtained through a variety of optimizations. The Offeror is strongly encouraged to optimize the source code in a variety of ways including (but not limited to):

- Using a different execution model (example: MPI+OpenMP/OpenACC/CUDA);
- Using vendor-specific hardware features to accelerate code;
- Running the benchmarks at a higher or lower concurrency, including MPI concurrency different from the Base Case; or
- Optimizing processor affinities and data layouts.

Aggressive code changes that enhance performance are also permitted as long as the full capabilities of the code are maintained, the code can still pass validation tests, and the underlying purpose of the benchmark is not compromised. Changes to the source code may be made so long as the following conditions are met:

- The rationale for and relative effect on performance of any optimization is described.
- Algorithms fundamental to the program are not replaced (since replacing algorithms may result in violations of correctness or program requirements or other chosen software decisions).
- All simulation parameters such as grid size, number of particles, etc., must not be changed.
- The optimized code execution must still result in correct numerical results.
- Any code optimizations must be made available to the general user community, either through a system library or a well-documented explanation of code improvements.
- Any library routine used must currently exist in an Offeror's supported set of general or scientific libraries, or must be in such a set when the system is delivered, and must not specialize or limit the applicability of the benchmark nor violate the measurement goals of the particular benchmark.
- Source preprocessors, execution profile feedback optimizers, etc., are allowed as long as they are, or will be, available and supported as part of the compilation environment for the delivered systems.
- Only publicly available and documented compiler switches shall be used.
- Finally, the same code optimizations must be made for all runs of a benchmark (meaning all sizes). For example, one set of code optimizations may not be made for the "small" case while a different set of optimizations is made for the "large" case.

Any specific code changes and the runtime configuration used must be clearly documented with a complete audit trail and all supporting documentation included in the submission. Trinity/NERSC-8 will be the final judge of whether optimizations will be acceptable.

SSP

Before reading the following description of the SSP and how it is used in this RFP, it is strongly recommended that the Offerer read and understand its prior definition, use, and history [1] as some terms in this brief description may otherwise be confusing. The SSP is a derived measure of computational capability relevant to achievable scientific work. It shall be used to validate the system and monitor delivered performance throughout the system lifecycle [1]. The SSP is derived from an application performance figure, P_i , expressed in units of TFLOP per second per node. Given a system configured with N nodes, the SSP is the geometric mean of P_i over all applications, multiplied by N . The floating-point operation counts used in calculating P_i for each of the applications have been pre-determined by Trinity/NERSC using either a hardware performance counter on a single reference system (NERSC's Hopper platform) or by actual counts within the code, and these values may *not* be altered. The floating-point operation counts are *not* measured on the Offeror's system; only the application run-time is to be measured. The reference TFLOP counts are to be found in the mini-applications section of the spreadsheet supplied with this RFP. For the purposes of the SSP calculation, the Offerer must use the run times and node counts of the large, optimized (MPI+X) problems described above.

As calculated in the manner given above, the SSP represents an "instantaneous" measure of computational capability as of the date the Offeror's application run times were measured. To represent the cumulative computational capability of a system over a specific period of time, the instantaneous SSP may be summed over that time period by multiplying the instantaneous value by the length of time. The result is expressed in units of rate-time; e.g., PFLOP/s-years .

In the SSP calculation there are two places where the number of nodes occurs: one is the number of nodes used to run the application benchmark and the other is the number of nodes with which the proposed system will be configured. Note that the former must be the physical number of nodes used to run the benchmark regardless of whether some processor elements in the nodes are left vacant (or if some are oversubscribed). ***The SSP is calculated based on the total number of physical nodes used, not (necessarily) the total concurrency.*** The Offeror should determine if the SSP increases or decreases when running on an unpacked mode.

References

1. <http://www.nersc.gov/research-and-development/performance-and-monitoring-tools/sustained-system-performance-ssp-benchmark>